

Claims

- [c1] 1. A method for regulating resource usage by a plurality of programs running on a plurality of machines, the method comprising:
 - providing a resource policy specifying allocation of resources amongst the plurality of programs;
 - determining resources available at the plurality of machines;
 - detecting requests for resources by each of the plurality of programs running on each of the plurality of machines;
 - periodically exchanging resource information amongst the plurality of machines, the resource information including requests for resources and resource availability at each of the plurality of machines; and
 - at each of the plurality of machines, allocating resources to each program based upon the resource policy and the resource information.
- [c2] 2. The method of claim 1, wherein said resources include communication resources.
- [c3] 3. The method of claim 2, wherein said communication resources include network bandwidth shared by the plu-

rality of machines.

- [c4] 4. The method of claim 1, wherein said resources include computing resources.
- [c5] 5. The method of claim 4, wherein said computing resources include processing resources available at the plurality of machines.
- [c6] 6. The method of claim 4, wherein said computing resources include memory resources available at the plurality of machines.
- [c7] 7. The method of claim 1, wherein the plurality of programs include an application that is running on a plurality of computers.
- [c8] 8. The method of claim 1, wherein the plurality of programs include an application that is running on a single computer.
- [c9] 9. The method of claim 1, wherein the resource policy includes a rule specifying a percentage of available resources to be allocated to a particular program.
- [c10] 10. The method of claim 1, wherein the resource policy includes a rule specifying a specific quantity of resources to be allocated to a particular program.

- [c11] 11. The method of claim 1, wherein the resource policy is user configurable.
- [c12] 12. The method of claim 1, wherein the resource policy specifies priorities for allocation of resources amongst the plurality of programs.
- [c13] 13. The method of claim 12, wherein the priorities for allocation of resources are automatically adjusted based on occurrence of particular events.
- [c14] 14. The method of claim 1, wherein said detecting step includes detecting each instance of a program running at each of the plurality of machines.
- [c15] 15. The method of claim 1, wherein said detecting step is performed with a frequency established by a user.
- [c16] 16. The method of claim 15, wherein a user can establish a frequency greater than once per second.
- [c17] 17. The method of claim 1, wherein said exchanging step is performed with a frequency established by a user.
- [c18] 18. The method of claim 17, wherein a user can establish a frequency greater than once per second.
- [c19] 19. The method of claim 1, wherein said exchanging step includes exchanging information based on changes

in resource availability since a prior exchange of information.

- [c20] 20. The method of claim 1, wherein said exchanging step includes exchanging information based on changes in requests for resources since a prior exchange of information.
- [c21] 21. The method of claim 1, wherein said exchanging step includes using a bandwidth-conserving protocol.
- [c22] 22. The method of claim 1, wherein said allocating step includes regulating usage of resources by each of the plurality of programs.
- [c23] 23. The method of claim 1, wherein said allocating step includes scheduling processing resources at each of the plurality of machines.
- [c24] 24. The method of claim 1, wherein said allocating step includes regulating the volume of communications sent by a particular program.
- [c25] 25. The method of claim 1, wherein said allocating step includes delaying the sending of a communication by a particular program.
- [c26] 26. The method of claim 1, further comprising:
collecting resource information regarding requests for

resources and resource availability; and generating resource utilization information for display to a user based upon the collected resource information.

- [c27] 27. The method of claim 26, further comprising: automatically suggesting modifications to the resource policy based, at least in part, upon the collected resource information.
- [c28] 28. A computer-readable medium having processor-executable instructions for performing the method of claim 1.
- [c29] 29. A downloadable set of processor-executable instructions for performing the method of claim 1.
- [c30] 30. A system for regulating utilization of computer resources of a plurality of computers, the system comprising:
 - a plurality of computers having resources to be regulated which are connected to each other through a network;
 - a monitoring module provided at each computer having resources to be regulated, for monitoring resource utilization and providing resource utilization information to each other connected computer having resources to be regulated;
 - a manager module providing rules governing utilization

of resources available on the plurality of computers and transferring said rules to the plurality of computers; and an enforcement module at each computer for which resources are to be regulated for regulating usage of resources based on said transferred rules and the resource utilization information received from other connected computers.

- [c31] 31. The system of claim 30, wherein said resources to be regulated include communication resources.
- [c32] 32. The system of claim 30, wherein said resources to be regulated include processing resources.
- [c33] 33. The system of claim 30, wherein said monitoring module at a given computer identifies at least one application running at the given computer.
- [c34] 34. The system of claim 33, wherein said monitoring module detects a request for resources by said at least one application.
- [c35] 35. The system of claim 33, wherein said monitoring module detects a request for network communication by said at least one application.
- [c36] 36. The system of claim 30, wherein said monitoring module at a given computer determines resources avail-

able at the given computer.

- [c37] 37. The system of claim 30, wherein said monitoring module at a given computer provides resource utilization information to each other connected computer at a fixed interval.
- [c38] 38. The system of claim 37, wherein said fixed interval is a sub-second interval.
- [c39] 39. The system of claim 37, wherein said fixed interval is configurable by a user.
- [c40] 40. The system of claim 30, wherein said monitoring module at a given computer provides resource utilization information to each other connected computer in response to particular events.
- [c41] 41. The system of claim 30, wherein said resource utilization information provided by said monitoring module includes information regarding requests for communication resources.
- [c42] 42. The system of claim 30, wherein said resource utilization information provided by said monitoring module includes information regarding requests for processing resources.
- [c43] 43. The system of claim 30, wherein said rules provided

by said manager module include a rule specifying a percentage of available resources to be allocated to a particular application.

- [c44] 44. The system of claim 30, wherein said rules provided by said manager module include a rule specifying a specific quantity of resources to be allocated to a particular application.
- [c45] 45. The system of claim 30, wherein said manager module permits a user to establish rules governing utilization of resources.
- [c46] 46. The system of claim 30, wherein said monitoring module uses a bandwidth-conserving protocol for providing resource utilization information.
- [c47] 47. The system of claim 30, wherein said enforcement module schedules processing resources at each of the plurality of computers based on said transferred rules and the resource utilization information.
- [c48] 48. The system of claim 30, wherein said enforcement module regulates the volume of communications sent by a particular application.
- [c49] 49. The system of claim 30, wherein said enforcement module regulates the frequency of communication by a

particular application.

- [c50] 50. The system of claim 30, further comprising:
a configuration module for a user to establish rules governing utilization of resources.
- [c51] 51. The system of claim 50, wherein said configuration module collects resource utilization information from the plurality of computers.
- [c52] 52. The system of claim 51, wherein said configuration module suggests rules governing utilization of resources based, at least in part, upon the collected resource utilization information.
- [c53] 53. The system of claim 51, wherein said configuration module displays the collected resource utilization information to a user.
- [c54] 54. A method for scheduling communications by a plurality of applications running on a plurality of computers connected to each other through a network, the method comprising:
providing a policy specifying priorities for scheduling communications by the plurality of applications;
periodically determining communication resources available at the plurality of computers;
at each of the plurality of computers, detecting requests

to communicate and identifying a particular application associated with each request; exchanging bandwidth information amongst the plurality of computers, the bandwidth information including applications making the requests to communicate and a measure of communications resources required to fulfill the requests; and at each of the plurality of computers, scheduling communications based upon the policy and the bandwidth information.

- [c55] 55. The method of claim 54, wherein said communications comprises incoming and outgoing network traffic.
- [c56] 56. The method of claim 54, wherein said communication resources include network bandwidth shared by the plurality of computers.
- [c57] 57. The method of claim 54, wherein said exchanging step occurs at a frequency established by a user.
- [c58] 58. The method of claim 57, wherein a user can establish a frequency greater than once per second.
- [c59] 59. The method of claim 54, wherein said scheduling step includes immediately transmitting all communications if the bandwidth information indicates communication traffic is light.

- [c60] 60. The method of claim 54, wherein said scheduling step includes delaying a portion of the communications if the bandwidth information indicates communication traffic is heavy.
- [c61] 61. The method of claim 60, wherein said scheduling step includes delaying transmission of communications by lower-priority applications.
- [c62] 62. The method of claim 60, wherein said scheduling step includes load balancing.
- [c63] 63. The method of claim 62, wherein said load balancing includes redirecting communications received at a first computer to a second computer.
- [c64] 64. A computer-readable medium having processor-executable instructions for performing the method of claim 54.
- [c65] 65. A downloadable set of processor-executable instructions for performing the method of claim 54.